

# Exhibit 4

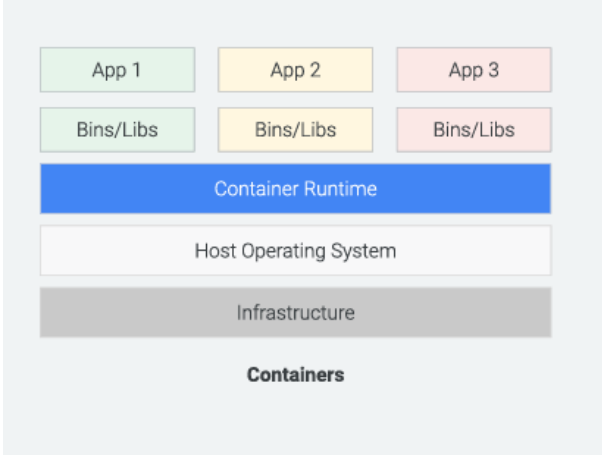
**U.S. Patent No. 7,784,058 (“’058 Patent”)**

Accused Instrumentalities: Google products and services using secure containerized applications, including without limitation Google Kubernetes Engine, Cloud Run, and Migrate to Containers, and all versions and variations thereof since the issuance of the asserted patent.

Each Accused Instrumentality infringes the claims in substantially the same way, and the evidence shown in this chart is similarly applicable to each Accused Instrumentality. Each claim limitation is literally infringed by each Accused Instrumentality. However, to the extent any claim limitation is not met literally, it is nonetheless met under the doctrine of equivalents because the differences between the claim limitation and each Accused Instrumentality would be insubstantial, and each Accused Instrumentality performs substantially the same function, in substantially the same way, to achieve the same result as the claimed invention. Notably, Defendant has not yet articulated which, if any, particular claim limitations it believes are not met by the Accused Instrumentalities.

**Claim 1**

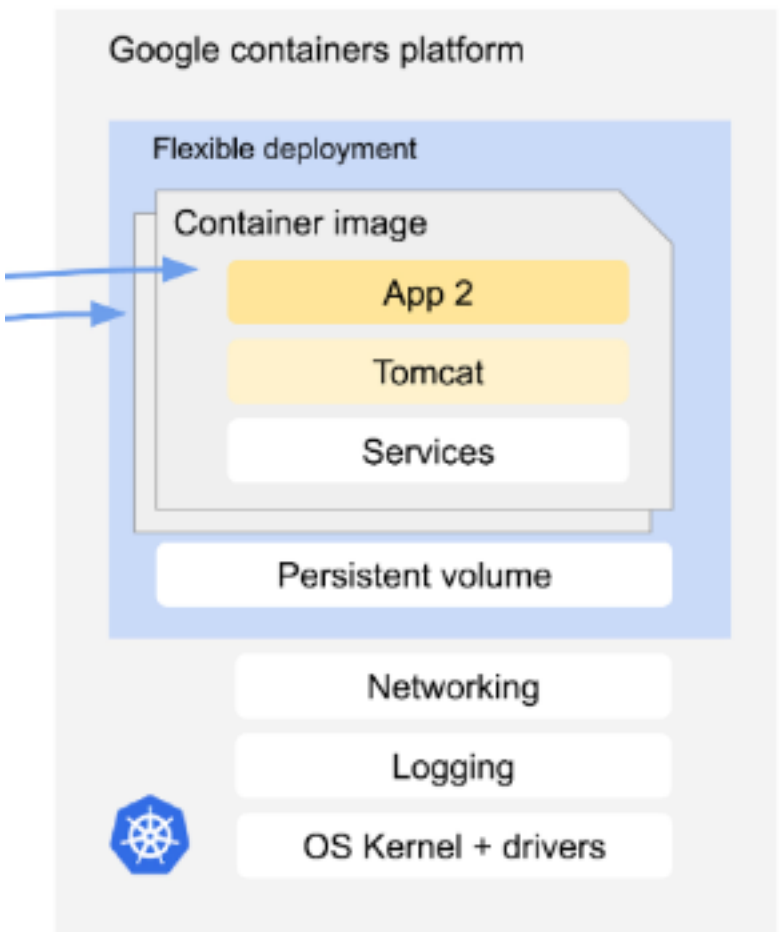
Claim 1	Accused Instrumentalities
<p>[1pre] 1. A computing system for executing a plurality of software applications comprising:</p>	<p>To the extent the preamble is limiting, each Accused Instrumentality comprises or constitutes a computing system for executing a plurality of software applications as claimed.</p> <p><i>See claim limitations below.</i></p> <p><i>See also, e.g.:</i></p> <p>Google Kubernetes Engine (GKE) clusters provide secured and managed Kubernetes services with autoscaling and multi-cluster support. GKE lets you deploy, manage, and scale containerized applications on Kubernetes, powered by Google Cloud.</p> <p><a href="https://cloud.google.com/migrate/containers/docs/getting-started">https://cloud.google.com/migrate/containers/docs/getting-started</a></p> <p>Use Migrate to Containers to modernize traditional applications away from virtual machine (VM) instances and into native containers that run on Google Kubernetes Engine (GKE), GKE Enterprise clusters, or Cloud Run platform. You can migrate workloads from VMs that run on VMware or Compute Engine, giving you the flexibility to containerize your existing workloads with ease. Migrate to Containers supports modernization of IBM WebSphere, JBoss, Apache, Tomcat, WordPress, Windows IIS applications, as well as containerisation of Linux-based applications.</p> <p><a href="https://cloud.google.com/migrate/containers/docs/getting-started">https://cloud.google.com/migrate/containers/docs/getting-started</a>.</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="688 277 1864 448">A container is a way of packaging a given application's code and dependencies so that the application will run easily in any computing environment. This solves the common problem of</p>  <p data-bbox="632 480 1230 932">The diagram illustrates the container architecture stack. At the top, three application boxes labeled 'App 1' (green), 'App 2' (yellow), and 'App 3' (pink) are shown. Below each app is a corresponding 'Bins/Libs' box of the same color. These are all supported by a blue 'Container Runtime' bar. Below the runtime is a light gray 'Host Operating System' bar, and at the bottom is a dark gray 'Infrastructure' bar. The entire stack is labeled 'Containers' at the bottom.</p> <p data-bbox="632 959 1629 992"><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 1	Accused Instrumentalities
	<p>Containers can run virtually anywhere, greatly easing development and deployment: on Linux, Windows, and Mac operating systems; on virtual machines or on physical servers; on a developer's machine or in data centers on-premises; and of course, in the public cloud.</p> <p>Containers are lightweight packages of your application code together with dependencies such as specific versions of programming language runtimes and libraries required to run your software services.</p> <p><a href="https://cloud.google.com/learn/what-are-containers">https://cloud.google.com/learn/what-are-containers</a></p>
[1a] a) a processor;	<p>Each Accused Instrumentality comprises a processor.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p>Containers virtualize CPU, memory, storage, and network resources at the operating system level, providing developers with a view of the OS logically isolated from other applications.</p> <p><a href="https://cloud.google.com/learn/what-are-containers">https://cloud.google.com/learn/what-are-containers</a></p> <ul style="list-style-type: none"> <li>• <b>Higher utilization and density</b>, leveraging automatic bin-packing and auto-scaling capabilities, Kubernetes places containers optimally in nodes based on required resources while scaling as needed, without impairing availability. In addition, unlike VMs, all containers on a single node share one copy of the operating system and don't each require their own OS image and vCPU, resulting in a much smaller memory footprint and CPU needs. This means more workloads running on fewer compute resources.</li> </ul> <p><a href="https://cloud.google.com/blog/products/containers-kubernetes/how-migrate-for-anthos-improves-vm-to-container-migration">https://cloud.google.com/blog/products/containers-kubernetes/how-migrate-for-anthos-improves-vm-to-container-migration</a></p> <p>Containers use specific features of the Linux kernel that "trick" individual applications into thinking they're in their own unique environment, even though multiple applications share the same host kernel. (If you're not familiar with the Linux kernel, it's a part of the operating system that communicates between processes--requests that do user tasks like opening a file, running a program-- and the hardware. It manages resources like memory and CPU to meet these requests).</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 1	Accused Instrumentalities
<p>[1b] b) an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor; and,</p>	<p>Each Accused Instrumentality comprises an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor.</p> <p><i>See, e.g.:</i></p> <ul style="list-style-type: none"> <li>• Containers are much more lightweight than VMs</li> <li>• Containers virtualize at the OS level while VMs virtualize at the hardware level</li> <li>• Containers share the OS kernel and use a fraction of the memory VMs require</li> </ul> <p><a href="https://cloud.google.com/learn/what-are-containers">https://cloud.google.com/learn/what-are-containers</a></p> <p><b>Kernel mode</b></p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p><a href="https://www.techtarget.com/searchdatacenter/definition/kernel">https://www.techtarget.com/searchdatacenter/definition/kernel</a></p>

Claim 1	Accused Instrumentalities
	 <p>The diagram illustrates the Google containers platform architecture. At the top is the 'Google containers platform' layer. Below it is the 'Flexible deployment' layer, which contains a 'Container image' box. This box is divided into three sections: 'App 2' (yellow), 'Tomcat' (yellow), and 'Services' (white). Two blue arrows point from the left towards the 'Container image' box. Below the 'Container image' box is a 'Persistent volume' box. Below the 'Flexible deployment' layer are three stacked boxes: 'Networking', 'Logging', and 'OS Kernel + drivers'. A Docker logo is positioned to the left of the 'OS Kernel + drivers' box.</p> <p><a href="https://cloud.google.com/blog/products/application-modernization/shift-your-apps-to-container-based-workloads-on-the-command-line">https://cloud.google.com/blog/products/application-modernization/shift-your-apps-to-container-based-workloads-on-the-command-line</a></p>

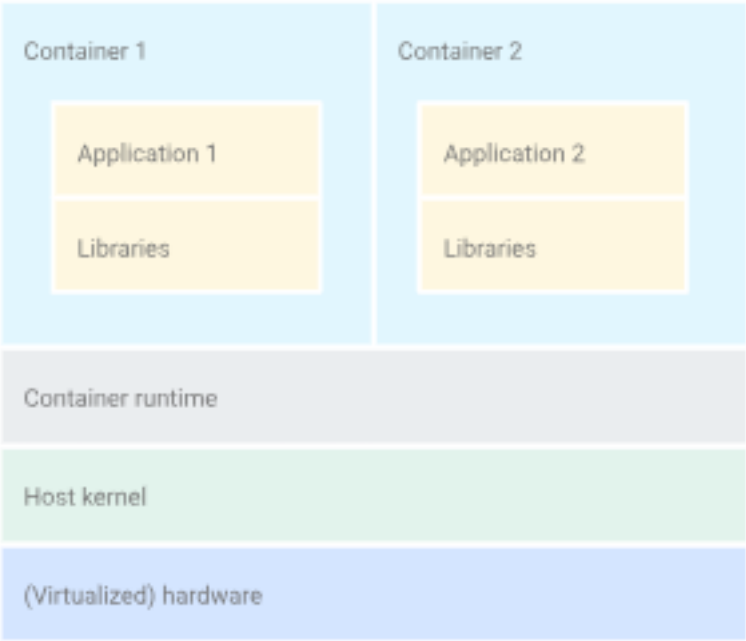
Claim 1	Accused Instrumentalities
	<p>The migration prerequisites are dependent on your specific migration environment. Confirm that your workloads' OS and source platform are compatible for migration by reviewing the prerequisites for your specific migration environment:</p> <p><a href="https://cloud.google.com/migrate/containers/docs/setting-up-overview">https://cloud.google.com/migrate/containers/docs/setting-up-overview</a></p> <p>Containers use specific features of the Linux kernel that “trick” individual applications into thinking they’re in their own unique environment, even though multiple applications share the same host kernel. (If you’re not familiar with the Linux kernel, it’s a part of the operating system that communicates between processes--requests that do user tasks like opening a file, running a program-- and the hardware. It manages resources like memory and CPU to meet these requests).</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p> <p>The <b>GNU C Library</b>, commonly known as <b>glibc</b>, is the <a href="#">GNU Project</a> implementation of the <a href="#">C standard library</a>. It is a wrapper around the system calls of the <a href="#">Linux kernel</a> for application use. Despite its name, it now also directly supports <a href="#">C++</a> (and, indirectly, other <a href="#">programming languages</a>). It was started in the 1980s by the <a href="#">Free Software Foundation</a> (FSF) for the <a href="#">GNU</a> operating system.</p> <p><a href="https://en.wikipedia.org/wiki/Glibc">https://en.wikipedia.org/wiki/Glibc</a></p>
[1c] c) a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode and	<p>Each Accused Instrumentality comprises a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode.</p> <p><i>See, e.g.:</i></p>



Claim 1	Accused Instrumentalities
	<p data-bbox="661 251 1228 690">A “container image” is your application and its dependencies, and uses a “base image” as the basis for the container image</p> <p data-bbox="661 755 1879 1177">The container image specifies the container’s file system. For example, if you’re running a Node.js application, the container image would contain your app, Node.js, and other dependencies like Linux system libraries (except the kernel). A container image usually extends a base operating system image, or <b>base image</b>. This base image is the basis of your container, so you’ll want to ensure that it’s properly patched and free from known vulnerabilities.</p> <p data-bbox="630 1218 1627 1250"><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

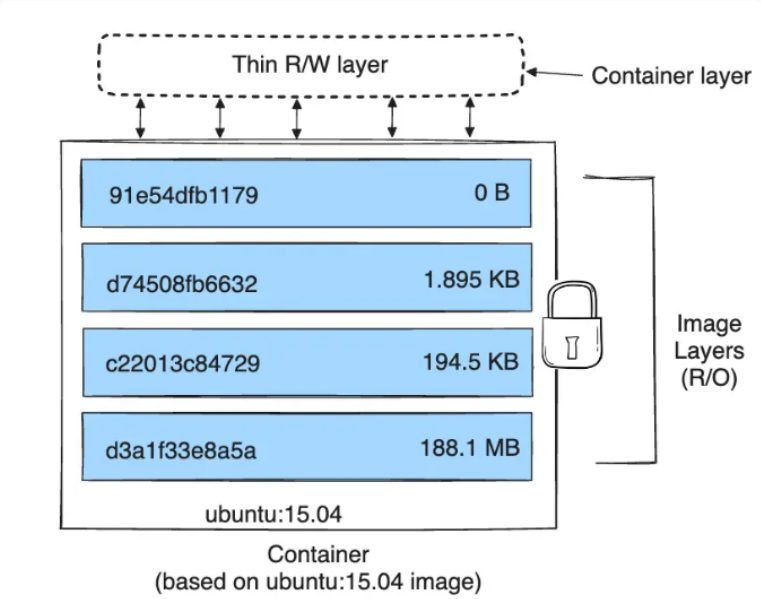
Claim 1	Accused Instrumentalities
	<p>A base image is the starting point for most container-based development workflows. Developers start with a base image and layer on top of it the necessary libraries, binaries, and configuration files used to run their application.</p> <p>Many base images are basic or minimal Linux distributions: Debian, Ubuntu, Red Hat Enterprise Linux (RHEL), Rocky Linux, or Alpine. Developers can consume these images directly from Docker Hub or other sources. There are official providers along with a wide variety of other downstream repackagers that layer software to meet customer needs.</p> <p>Google maintains base images for building its own applications. These images are built from the same source that Docker Hub uses. Therefore, they match the images you would get from Docker Hub.</p> <p><a href="https://cloud.google.com/software-supply-chain-security/docs/base-images">https://cloud.google.com/software-supply-chain-security/docs/base-images</a></p> <p>The <a href="#">preconfigured base images</a> provided by Cloud Workstations contain only a minimal environment with IDE, basic Linux terminal and language tools and a <code>sshd</code> server. To expedite the environment setup of specific development use cases, you can create custom container images that extend these base images to pre-install tools and dependencies and that run automation scripts.</p> <p>For custom container images, we recommend setting up a pipeline to automatically rebuild these images when the Cloud Workstations base image is updated, in addition to running a container scanning tool such as <a href="#">Artifact Analysis</a> to inspect any additional dependencies you added. You're responsible for maintaining and updating custom packages and dependencies added to custom images.</p> <p><a href="https://cloud.google.com/workstations/docs/customize-container-images">https://cloud.google.com/workstations/docs/customize-container-images</a></p>

Claim 1	Accused Instrumentalities
	<p>A container is a way of packaging a given application's code and dependencies so that the application will run easily in any computing environment. This solves the common problem of</p> <p>Containers solve the portability problem by isolating the application and its dependencies so they can be moved seamlessly between machines. A process running in a container lives isolated from the underlying environment. You control what it can see and what resources it can access. This helps you use resources more efficiently and not worry about the underlying infrastructure.</p> <p>One of the primary reasons to adopt containers is for your applications to be decoupled from the underlying environment and support higher resource utilization by "bin packing" multiple workloads onto each server. As such, the architecture of containers means that they're deployed with multiple containers sharing the same kernel.</p> <p>The core components of the Linux kernel that are used for containers are <b>cgroups</b> — control groups, which define the resources like CPU and memory which are available to a given process — and <b>namespaces</b>, which are a way of separating processes by restricting what each process can see, so that system resources "appear" isolated to the process.</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 1	Accused Instrumentalities
	 <p>The diagram illustrates a container architecture stack. At the top, two light blue boxes represent 'Container 1' and 'Container 2'. Inside 'Container 1' are two yellow boxes: 'Application 1' and 'Libraries'. Similarly, 'Container 2' contains 'Application 2' and 'Libraries'. Below the containers is a grey box for 'Container runtime', followed by a green box for 'Host kernel', and a blue box for '(Virtualized) hardware' at the base.</p> <p><a href="https://cloud.google.com/architecture/best-practices-for-operating-containers">https://cloud.google.com/architecture/best-practices-for-operating-containers</a></p> <p>For example, Migrate to Containers automatically generates a container image, a Dockerfile for day-2 image updates and application revisions, Kubernetes deployment YAMLs and (where relevant) a persistent data volume onto which the application data files and persistent state are copied. This automated, intelligent extraction is</p> <p><a href="https://cloud.google.com/blog/products/containers-kubernetes/how-migrate-for-anthos-improves-vm-to-container-migration">https://cloud.google.com/blog/products/containers-kubernetes/how-migrate-for-anthos-improves-vm-to-container-migration</a></p>

Claim 1	Accused Instrumentalities
	<p>Containers are lightweight packages of your application code together with dependencies such as specific versions of programming language runtimes and libraries required to run your software services.</p> <p><a href="https://cloud.google.com/learn/what-are-containers">https://cloud.google.com/learn/what-are-containers</a></p> <h2>About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2>Storage drivers versus Docker volumes</h2> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a href="#">volumes section</a> to learn how to use volumes to persist data and improve performance.</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="657 245 1083 302">Images and layers</h2> <p data-bbox="657 337 1822 415">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="674 483 1453 797"> # syntax=docker/dockerfile:1  FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="657 862 1900 1170">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="632 1192 1226 1224"><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>




Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the <a href="#">Best practices for writing Dockerfiles</a> and <a href="#">use multi-stage builds</a> sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer architecture of a Docker container. It shows a stack of four image layers, each represented by a blue rectangle with a unique ID and size:     <ul style="list-style-type: none"> <li>Top layer: ID <code>91e54dfb1179</code>, size <code>0 B</code></li> <li>Second layer: ID <code>d74508fb6632</code>, size <code>1.895 KB</code></li> <li>Third layer: ID <code>c22013c84729</code>, size <code>194.5 KB</code></li> <li>Bottom layer: ID <code>d3a1f33e8a5a</code>, size <code>188.1 MB</code></li> </ul>     The bottom layer is labeled <code>ubuntu:15.04</code>. To the right of the stack is a padlock icon and the text "Image Layers (R/O)", indicating these are read-only. Above the stack is a dashed box labeled "Thin R/W layer", with an arrow pointing to it from the text "Container layer". Vertical double-headed arrows connect the thin R/W layer to each of the four image layers below it. Below the entire stack is the text "Container (based on ubuntu:15.04 image)".   </p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 256 919 321">Volumes</h2> <p data-bbox="653 375 1906 505">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While <a href="#">bind mounts</a> are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="634 526 1308 558"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p> <h2 data-bbox="653 610 1226 659">Container environment</h2> <p data-bbox="653 696 1474 764">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="695 802 1451 964" style="list-style-type: none"><li>• A filesystem, which is a combination of an <a href="#">image</a> and one or more <a href="#">volumes</a>.</li><li>• Information about the Container itself.</li><li>• Information about other objects in the cluster.</li></ul> <p data-bbox="634 997 1528 1029"><a href="https://kubernetes.io/docs/concepts/containers/container-environment/">https://kubernetes.io/docs/concepts/containers/container-environment/</a></p>



Claim 1	Accused Instrumentalities
	<h2 data-bbox="659 250 877 315">Images</h2> <p data-bbox="659 347 1522 500">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="659 537 1528 607">You typically create a container image of your application and push it to a registry before referring to it in a <u>Pod</u>.</p> <p data-bbox="634 634 1329 667"><a href="https://kubernetes.io/docs/concepts/containers/images/">https://kubernetes.io/docs/concepts/containers/images/</a></p> <h2 data-bbox="653 711 919 769">Volumes</h2> <p data-bbox="653 808 1482 878">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers.</p> <p data-bbox="653 889 1430 922">One problem occurs when a container crashes or is stopped.</p> <p data-bbox="653 933 1528 1252">Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a <code>Pod</code> and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes <u>volume</u> abstraction solves both of these problems. Familiarity with <code>Pods</code> is suggested.</p> <p data-bbox="634 1279 1308 1312"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p>

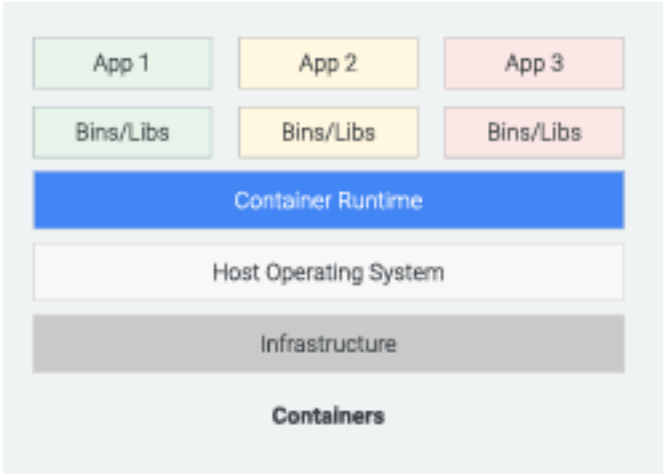
Claim 1	Accused Instrumentalities
	<div data-bbox="667 256 1297 316"><h2>Open Container Initiative</h2></div> <div data-bbox="667 378 1184 423"><h3>Image Format Specification</h3></div> <div data-bbox="667 472 1892 545"><p>This specification defines an OCI Image, consisting of an <a href="#">image manifest</a>, an <a href="#">image index</a> (optional), a set of <a href="#">filesystem layers</a>, and a <a href="#">configuration</a>.</p></div> <div data-bbox="667 581 1900 654"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="632 683 1478 751"><p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a></p></div>

Claim 1	Accused Instrumentalities
	<p data-bbox="653 250 831 289"><b>Overview</b></p> <p data-bbox="653 345 1892 586">At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more <a href="#">filesystem layer changeset</a> archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p> <div data-bbox="653 630 1892 1008"> <pre data-bbox="674 764 982 857"> public class HelloWorld {     public static void main(String[] args) {         System.out.println("Hello, World");     } } </pre> <p data-bbox="1079 634 1310 959">    /bin/java  /opt/app.jar  /lib/libc </p> <p data-bbox="1171 979 1234 1008">layer</p> <p data-bbox="1339 813 1367 841">+</p> <p data-bbox="1367 634 1612 959">    {   "manifests": {     "platform": {       "os": "linux",       ...     }   } } </p> <p data-bbox="1430 979 1577 1008">image index</p> <p data-bbox="1633 813 1661 841">+</p> <p data-bbox="1661 634 1906 959">    {   ...   "config": {     "Cmd": [       "java", "-jar",       "app.jar"     ],     ...   } } </p> <p data-bbox="1759 979 1829 1008">config</p> </div> <p data-bbox="632 1036 1478 1105"> <a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a> </p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 245 1297 305">OCI Image Configuration</h2> <p data-bbox="653 358 1913 521">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in <a href="#">Layers</a>.</p> <p data-bbox="653 558 1661 591">This section defines the <code>application/vnd.oci.image.config.v1+json</code> <a href="#">media type</a>.</p> <p data-bbox="632 623 1507 695"><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

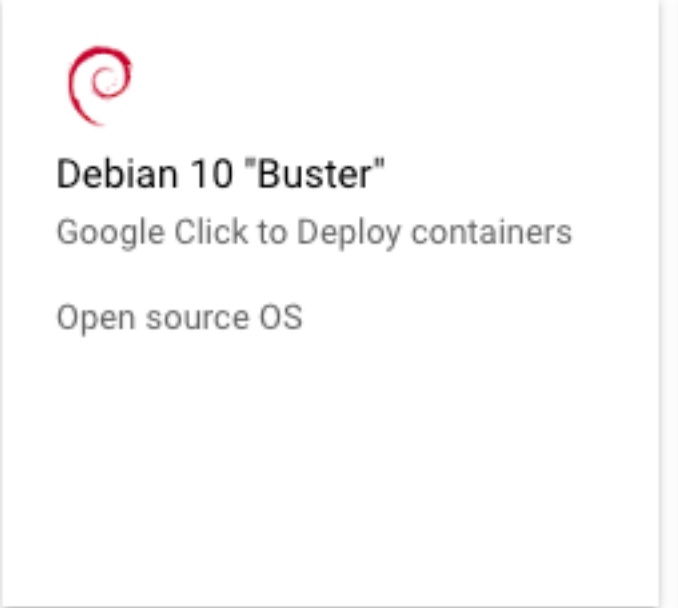
Claim 1	Accused Instrumentalities
	<p><b>Layer</b></p> <ul style="list-style-type: none"> <li>• Image filesystems are composed of <i>layers</i>.</li> <li>• Each layer represents a set of filesystem changes in a tar-based <a href="#">layer format</a>, recording files to be added, changed, or deleted relative to its parent layer.</li> <li>• Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer.</li> <li>• Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem.</li> </ul> <p><b>Image JSON</b></p> <ul style="list-style-type: none"> <li>• Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes.</li> <li>• The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers.</li> <li>• This JSON is considered to be immutable, because changing it would change the computed <a href="#">ImageID</a>.</li> <li>• Changing it means creating a new derived image, instead of changing the existing image.</li> </ul> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>
[1d] i) wherein some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible	In each Accused Instrumentality, some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when


Claim 1	Accused Instrumentalities
<p>to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications,</p>	<p>one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications.</p> <p>For example, a base image serves as a self-contained unit that encompasses all the necessary components for an application to run, including the application code, runtime environment, system tools, and dependencies (i.e., SLCSEs). The images are based on existing Linux distributions, such as Debian and Ubuntu, including essential system elements (i.e., functional replicas of OSCSEs). Each container image is based on a specific base image, which contains the application code, and dependencies, including system libraries or shared library critical system elements (SLCSEs). When the container runs the image, it creates a runtime instance of that container image.</p> <p><i>See, e.g.:</i></p> <p>Many base images are basic or minimal Linux distributions: Debian, Ubuntu, Red Hat Enterprise Linux (RHEL), Rocky Linux, or Alpine. Developers can consume these images directly from Docker Hub or other sources. There are official providers along with a wide variety of other downstream repackagers that layer software to meet customer needs.</p> <p><a href="https://cloud.google.com/software-supply-chain-security/docs/base-images">https://cloud.google.com/software-supply-chain-security/docs/base-images</a></p> <p>A container is a way of packaging a given application's code and dependencies so that the application will run easily in any computing environment. This solves the common problem of</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="646 245 1012 521">A “container image” is your application and its dependencies, and uses a “base image” as the basis for the container image</p>  <p>The diagram illustrates the container architecture stack. At the top, three application boxes labeled 'App 1' (green), 'App 2' (yellow), and 'App 3' (pink) are shown. Below each application is a corresponding 'Bins/Libs' box in the same color. These applications and their dependencies sit on a blue 'Container Runtime' bar. This bar is supported by a light gray 'Host Operating System' bar, which in turn sits on a dark gray 'Infrastructure' bar. The entire stack is labeled 'Containers' at the bottom.</p>

Claim 1	Accused Instrumentalities																								
	<p>The container image specifies the container's file system. For example, if you're running a Node.js application, the container image would contain your app, Node.js, and other dependencies like Linux system libraries (except the kernel). A container image usually extends a base operating system image, or <b>base image</b>. This base image is the basis of your container, so you'll want to ensure that it's properly patched and free from known vulnerabilities.</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p> <table><tr><th>OS</th><th>Repository path</th><th>Google Cloud Marketplace listing</th></tr><tr><td>Debian 10 "Buster"</td><td><a href="https://marketplace.gcr.io/google/debian10">marketplace.gcr.io/google/debian10</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr><tr><td>Debian 11 "Bullseye"</td><td><a href="https://marketplace.gcr.io/google/debian11">marketplace.gcr.io/google/debian11</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr><tr><td>Debian 12 "Bookworm"</td><td><a href="https://marketplace.gcr.io/google/debian12">marketplace.gcr.io/google/debian12</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr><tr><td>Rocky Linux 8</td><td><a href="https://marketplace.gcr.io/google/rockylinux8">marketplace.gcr.io/google/rockylinux8</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr><tr><td>Rocky Linux 9</td><td><a href="https://marketplace.gcr.io/google/rockylinux9">marketplace.gcr.io/google/rockylinux9</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr><tr><td>Ubuntu 20.04</td><td><a href="https://marketplace.gcr.io/google/ubuntu2004">marketplace.gcr.io/google/ubuntu2004</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr><tr><td>Ubuntu 22.04</td><td><a href="https://marketplace.gcr.io/google/ubuntu2204">marketplace.gcr.io/google/ubuntu2204</a></td><td><a href="#">Google Cloud Marketplace</a></td></tr></table> <p><a href="https://cloud.google.com/software-supply-chain-security/docs/base-images">https://cloud.google.com/software-supply-chain-security/docs/base-images</a></p>	OS	Repository path	Google Cloud Marketplace listing	Debian 10 "Buster"	<a href="https://marketplace.gcr.io/google/debian10">marketplace.gcr.io/google/debian10</a>	<a href="#">Google Cloud Marketplace</a>	Debian 11 "Bullseye"	<a href="https://marketplace.gcr.io/google/debian11">marketplace.gcr.io/google/debian11</a>	<a href="#">Google Cloud Marketplace</a>	Debian 12 "Bookworm"	<a href="https://marketplace.gcr.io/google/debian12">marketplace.gcr.io/google/debian12</a>	<a href="#">Google Cloud Marketplace</a>	Rocky Linux 8	<a href="https://marketplace.gcr.io/google/rockylinux8">marketplace.gcr.io/google/rockylinux8</a>	<a href="#">Google Cloud Marketplace</a>	Rocky Linux 9	<a href="https://marketplace.gcr.io/google/rockylinux9">marketplace.gcr.io/google/rockylinux9</a>	<a href="#">Google Cloud Marketplace</a>	Ubuntu 20.04	<a href="https://marketplace.gcr.io/google/ubuntu2004">marketplace.gcr.io/google/ubuntu2004</a>	<a href="#">Google Cloud Marketplace</a>	Ubuntu 22.04	<a href="https://marketplace.gcr.io/google/ubuntu2204">marketplace.gcr.io/google/ubuntu2204</a>	<a href="#">Google Cloud Marketplace</a>
OS	Repository path	Google Cloud Marketplace listing																							
Debian 10 "Buster"	<a href="https://marketplace.gcr.io/google/debian10">marketplace.gcr.io/google/debian10</a>	<a href="#">Google Cloud Marketplace</a>																							
Debian 11 "Bullseye"	<a href="https://marketplace.gcr.io/google/debian11">marketplace.gcr.io/google/debian11</a>	<a href="#">Google Cloud Marketplace</a>																							
Debian 12 "Bookworm"	<a href="https://marketplace.gcr.io/google/debian12">marketplace.gcr.io/google/debian12</a>	<a href="#">Google Cloud Marketplace</a>																							
Rocky Linux 8	<a href="https://marketplace.gcr.io/google/rockylinux8">marketplace.gcr.io/google/rockylinux8</a>	<a href="#">Google Cloud Marketplace</a>																							
Rocky Linux 9	<a href="https://marketplace.gcr.io/google/rockylinux9">marketplace.gcr.io/google/rockylinux9</a>	<a href="#">Google Cloud Marketplace</a>																							
Ubuntu 20.04	<a href="https://marketplace.gcr.io/google/ubuntu2004">marketplace.gcr.io/google/ubuntu2004</a>	<a href="#">Google Cloud Marketplace</a>																							
Ubuntu 22.04	<a href="https://marketplace.gcr.io/google/ubuntu2204">marketplace.gcr.io/google/ubuntu2204</a>	<a href="#">Google Cloud Marketplace</a>																							

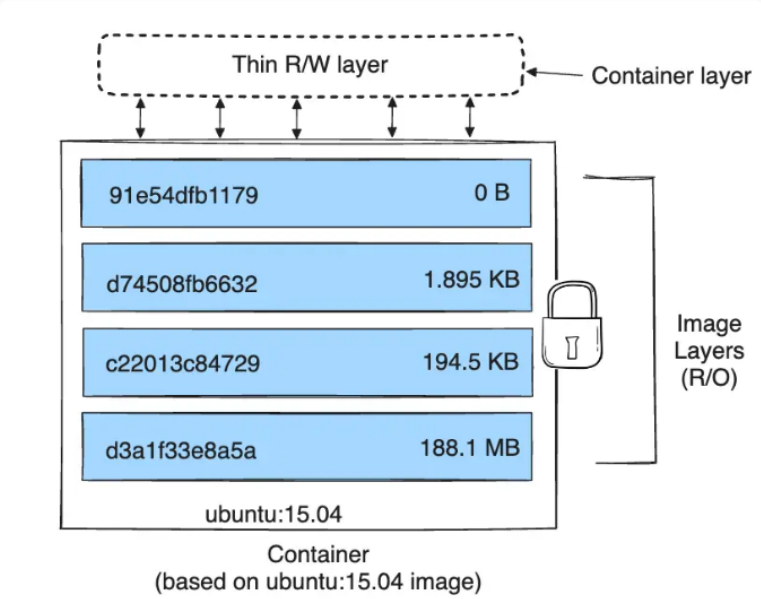


Claim 1	Accused Instrumentalities
	

Claim 1	Accused Instrumentalities
	<div data-bbox="682 272 766 354"></div> <div data-bbox="672 375 928 423">Ubuntu 20.04</div> <div data-bbox="672 436 1205 483">Google Click to Deploy containers</div> <div data-bbox="625 862 1717 904"><a href="https://console.cloud.google.com/marketplace/browse?filter=solution-type:container">https://console.cloud.google.com/marketplace/browse?filter=solution-type:container</a></div>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="646 248 1272 316">About storage drivers</h2> <p data-bbox="646 363 1871 488">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="646 557 1564 613">Storage drivers versus Docker volumes</h2> <p data-bbox="646 651 1915 917">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="646 967 1904 1092">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1339 1015 1535 1040" href="#">volumes section</a> to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="636 1123 1224 1154"><a data-bbox="636 1123 1224 1154" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="657 245 1081 302">Images and layers</h2> <p data-bbox="657 337 1822 415">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="674 483 1451 797"> # syntax=docker/dockerfile:1  FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="657 862 1902 1170">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="632 1192 1226 1224"><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>




Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the <a href="#">Best practices for writing Dockerfiles</a> and <a href="#">use multi-stage builds</a> sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer architecture of a Docker container. At the top, a dashed box labeled "Thin R/W layer" (also indicated by an arrow as the "Container layer") represents the writable layer for the container. Below this is a stack of four solid blue boxes representing the "Image Layers (R/O)". These layers are labeled with their commit IDs and sizes: "91e54dfb1179" (0 B), "d74508fb6632" (1.895 KB), "c22013c84729" (194.5 KB), and "d3a1f33e8a5a" (188.1 MB). A padlock icon is shown next to the stack, indicating that these layers are read-only. The entire stack is labeled "ubuntu:15.04". Below the stack, the text "Container (based on ubuntu:15.04 image)" is displayed.</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 256 919 321">Volumes</h2> <p data-bbox="653 375 1906 505">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While <a href="#">bind mounts</a> are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="634 526 1308 558"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p> <h2 data-bbox="653 610 1226 659">Container environment</h2> <p data-bbox="653 696 1474 764">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="695 802 1451 964" style="list-style-type: none"><li>• A filesystem, which is a combination of an <a href="#">image</a> and one or more <a href="#">volumes</a>.</li><li>• Information about the Container itself.</li><li>• Information about other objects in the cluster.</li></ul> <p data-bbox="634 997 1528 1029"><a href="https://kubernetes.io/docs/concepts/containers/container-environment/">https://kubernetes.io/docs/concepts/containers/container-environment/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="659 250 877 315">Images</h2> <p data-bbox="659 347 1522 500">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="659 537 1528 607">You typically create a container image of your application and push it to a registry before referring to it in a <u>Pod</u>.</p> <p data-bbox="634 634 1329 667"><a href="https://kubernetes.io/docs/concepts/containers/images/">https://kubernetes.io/docs/concepts/containers/images/</a></p> <h2 data-bbox="653 711 919 769">Volumes</h2> <p data-bbox="653 808 1482 878">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers.</p> <p data-bbox="653 889 1430 922">One problem occurs when a container crashes or is stopped.</p> <p data-bbox="653 933 1528 1252">Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a <code>Pod</code> and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes <u>volume</u> abstraction solves both of these problems. Familiarity with <code>Pods</code> is suggested.</p> <p data-bbox="634 1279 1308 1312"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p>

Claim 1	Accused Instrumentalities
	<div data-bbox="667 256 1297 316"><h2>Open Container Initiative</h2></div> <div data-bbox="667 378 1184 423"><h3>Image Format Specification</h3></div> <div data-bbox="667 472 1892 545"><p>This specification defines an OCI Image, consisting of an <a href="#">image manifest</a>, an <a href="#">image index</a> (optional), a set of <a href="#">filesystem layers</a>, and a <a href="#">configuration</a>.</p></div> <div data-bbox="667 581 1900 654"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="632 683 1478 751"><p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a></p></div>

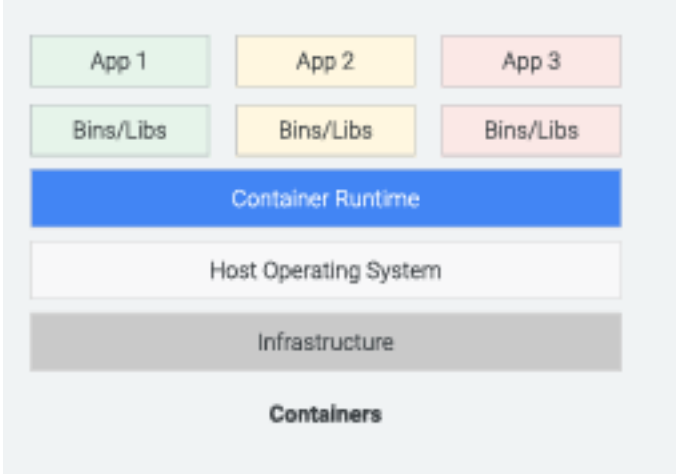


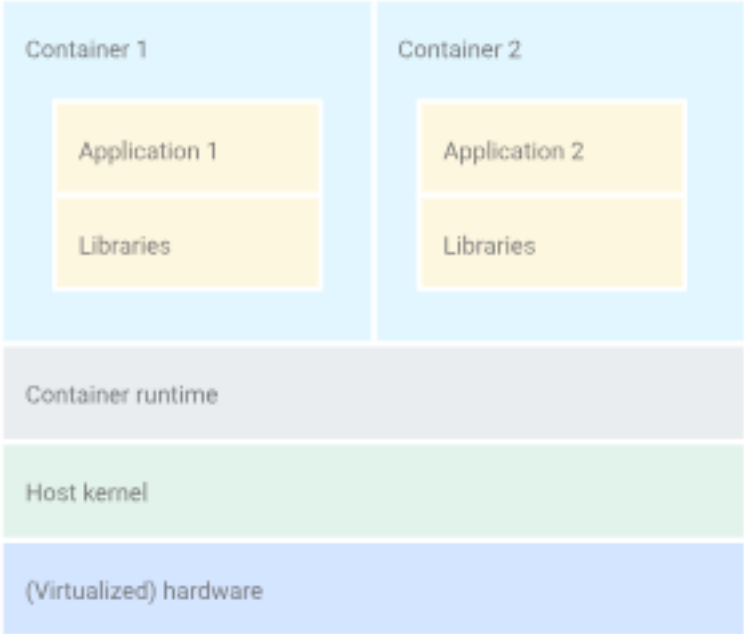
Claim 1	Accused Instrumentalities
	<p data-bbox="648 250 831 289"><b>Overview</b></p> <p data-bbox="648 345 1892 586">At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more <a href="#">filesystem layer changeset</a> archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p> <div data-bbox="667 626 1906 1008"> <pre data-bbox="667 764 982 857"> public class HelloWorld {     public static void main(String[] args) {         System.out.println("Hello, World");     } } </pre> <p data-bbox="1077 634 1312 959">    /bin/java  /opt/app.jar  /lib/libc </p> <p data-bbox="1171 979 1234 1008">layer</p> <p data-bbox="1339 808 1367 829">+</p> <p data-bbox="1381 634 1617 959">    {   "manifests": {     "platform": {       "os": "linux",       ...     }   } } </p> <p data-bbox="1430 979 1577 1008">image index</p> <p data-bbox="1633 808 1661 829">+</p> <p data-bbox="1675 634 1906 959">    {   ...   "config": {     "Cmd": [       "java", "-jar",       "app.jar"     ],     ...   } } </p> <p data-bbox="1759 979 1829 1008">config</p> </div> <p data-bbox="632 1036 1478 1105"> <a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a> </p>


Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 245 1297 305">OCI Image Configuration</h2> <p data-bbox="653 358 1913 521">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in <a href="#">Layers</a>.</p> <p data-bbox="653 558 1661 591">This section defines the <code>application/vnd.oci.image.config.v1+json</code> <a href="#">media type</a>.</p> <p data-bbox="632 623 1507 695"><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Layer</b></p> <ul style="list-style-type: none"> <li>• Image filesystems are composed of <i>layers</i>.</li> <li>• Each layer represents a set of filesystem changes in a tar-based <a href="#">layer format</a>, recording files to be added, changed, or deleted relative to its parent layer.</li> <li>• Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer.</li> <li>• Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem.</li> </ul> <p><b>Image JSON</b></p> <ul style="list-style-type: none"> <li>• Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes.</li> <li>• The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers.</li> <li>• This JSON is considered to be immutable, because changing it would change the computed <a href="#">ImageID</a>.</li> <li>• Changing it means creating a new derived image, instead of changing the existing image.</li> </ul> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>
[1e] ii) wherein an instance of a SLCSE provided to at least a first of the plurality of software applications from the	In each Accused Instrumentality, an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the

Claim 1	Accused Instrumentalities
<p>shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function, and</p>	<p>operating system have use of a unique instance of a corresponding critical system element for performing same function.</p> <p>When a Docker or Kubernetes image is used to create a container, it creates a separate and isolated instance of a runtime environment which is independent of other containers running on the same host. Each container has its own instance of base images and its own data. The containers run in isolation, ensuring that the SLCSEs stored in the shared library are accessible to the software applications running in their respective containers. The image includes essential system files, libraries, and dependencies required to run the software application within the container. The containers can share common dependencies and components using layered images. This means that multiple containers utilize the same base image to create an instance. When an instance of SLCSE is provided from the base image (i.e., from the shared library) to an individual container including application software, it operates in isolation and runs its own instance of the software application without sharing resources or critical system elements with other containers. This ensures that each container has its own isolated context. Docker or Kubernetes containers can share common dependencies and components using layered images. This means that multiple containers can utilize the same base image. Therefore, each container, containing the application software running under the operating system, utilizes a unique instance of the corresponding critical system element to execute the respective application software for performing a same or a different function.</p> <p><i>See, e.g.:</i></p> <p>A container is a way of packaging a given application's code and dependencies so that the application will run easily in any computing environment. This solves the common problem of</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 1	Accused Instrumentalities
	<p>Containers solve the portability problem by isolating the application and its dependencies so they can be moved seamlessly between machines. A process running in a container lives isolated from the underlying environment. You control what it can see and what resources it can access. This helps you use resources more efficiently and not worry about the underlying infrastructure.</p> <p>The container image specifies the container's file system. For example, if you're running a Node.js application, the container image would contain your app, Node.js, and other dependencies like Linux system libraries (except the kernel). A container image usually extends a base operating system image, or <b>base image</b>. This base image is the basis of your container, so you'll want to ensure that it's properly patched and free from known vulnerabilities.</p>  <p>The diagram illustrates the container architecture stack. At the top, three application boxes labeled 'App 1' (green), 'App 2' (yellow), and 'App 3' (pink) are shown. Below each app is a corresponding 'Bins/Libs' box in the same color. These are all contained within a blue box labeled 'Container Runtime'. Below the runtime is a white box labeled 'Host Operating System', and at the bottom is a grey box labeled 'Infrastructure'. The entire stack is labeled 'Containers' at the bottom.</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 1	Accused Instrumentalities
	 <p>The diagram illustrates a container architecture stack. At the top, two containers are shown side-by-side: 'Container 1' and 'Container 2'. Inside 'Container 1' are 'Application 1' and 'Libraries'. Inside 'Container 2' are 'Application 2' and 'Libraries'. Below the containers is a grey bar labeled 'Container runtime'. Below that is a green bar labeled 'Host kernel'. At the bottom is a blue bar labeled '(Virtualized) hardware'.</p> <p><a href="https://cloud.google.com/architecture/best-practices-for-operating-containers">https://cloud.google.com/architecture/best-practices-for-operating-containers</a></p>

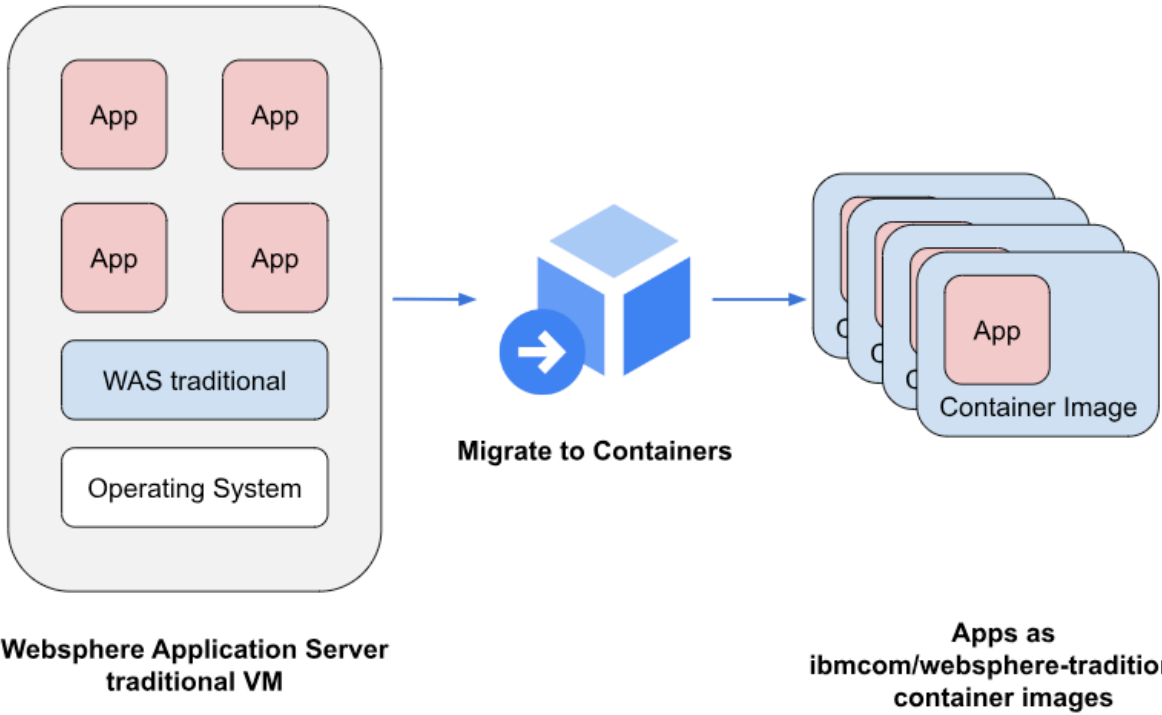


Claim 1	Accused Instrumentalities
	<div data-bbox="699 272 1824 950"> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>Dockerfile App 1</p> <div style="background-color: #f8d7da; padding: 5px; margin-bottom: 5px;">FROM node:19.7.0</div> <div style="background-color: #d4edda; padding: 5px; margin-bottom: 5px;">ADD src_app1 /src/</div> <div style="background-color: #d4edda; padding: 5px;">RUN cd /src &amp;&amp; \npm install</div> </div> <div style="text-align: center;"> <p>Dockerfile App 2</p> <div style="background-color: #f8d7da; padding: 5px; margin-bottom: 5px;">FROM node:19.7.0</div> <div style="background-color: #d4edda; padding: 5px; margin-bottom: 5px;">ADD src_app2 /src/</div> <div style="background-color: #d4edda; padding: 5px;">RUN cd /src &amp;&amp; \npm install</div> </div> </div> <div style="margin-top: 10px;"> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="width: 15px; height: 15px; background-color: #f8d7da; margin-right: 5px;"></div> <span>Common layers, downloaded only once</span> </div> <div style="display: flex; align-items: center;"> <div style="width: 15px; height: 15px; background-color: #d4edda; margin-right: 5px;"></div> <span>Layers unique to each image</span> </div> </div> </div> <p data-bbox="632 1015 1680 1052"><a href="https://cloud.google.com/architecture/best-practices-for-building-containers">https://cloud.google.com/architecture/best-practices-for-building-containers</a></p> <p data-bbox="632 1089 1917 1198">One method of packaging an application into a container is with the use of a Dockerfile. The Dockerfile is similar to a script which instructs the daemon on how to assemble the container image. See the <a href="#">Dockerfile reference documentation</a> for more information.</p> <p data-bbox="632 1224 1917 1333">Using the Dockerfile method to build a container requires direct knowledge about the application in order to assemble the container. The first step to creating a Dockerfile is selecting an image that will be used as the basis of your image. This image should be a parent or base image maintained and published by a trusted source, usually your company.</p> <p data-bbox="632 1352 1768 1390"><a href="https://codelabs.developers.google.com/developing-containers-with-dockerfiles#2">https://codelabs.developers.google.com/developing-containers-with-dockerfiles#2</a></p>



Claim 1	Accused Instrumentalities
<p>[1f] iii) wherein a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p>	<p>In each Accused Instrumentality, a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p> <p>For example, in Docker or Kubernetes containers, each container operates independently, and a base image includes essential system files, libraries, and dependencies (i.e., SLCSEs) required to run the software application within the container. Based on information and belief, each element, such as system files, libraries, and dependencies (i.e., SLCSE) is associated with an execution of a predetermined function related to the application. When an image is used to create a container in the Accused Instrumentality, an instance of the SLCSE is provided to a software application. Therefore, different instances of the SLCSE are provided to different applications for performing either a same or a different function, simultaneously.</p> <p><i>See, e.g.:</i></p> <p>Containers solve the portability problem by isolating the application and its dependencies so they can be moved seamlessly between machines. A process running in a container lives isolated from the underlying environment. You control what it can see and what resources it can access. This helps you use resources more efficiently and not worry about the underlying infrastructure.</p> <p>The container image specifies the container's file system. For example, if you're running a Node.js application, the container image would contain your app, Node.js, and other dependencies like Linux system libraries (except the kernel). A container image usually extends a base operating system image, or <b>base image</b>. This base image is the basis of your container, so you'll want to ensure that it's properly patched and free from known vulnerabilities.</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 1	Accused Instrumentalities
	<div data-bbox="699 272 1824 950"> <div> <div>Dockerfile App 1</div> <div>FROM node:19.7.0</div> <div>ADD src_app1 /src/</div> <div>RUN cd /src &amp;&amp; \npm install</div> </div> <div> <div>Dockerfile App 2</div> <div>FROM node:19.7.0</div> <div>ADD src_app2 /src/</div> <div>RUN cd /src &amp;&amp; \npm install</div> </div> <div> <div>Common layers, downloaded only once</div> <div>Layers unique to each image</div> </div> </div> <p data-bbox="632 1015 1680 1052"><a href="https://cloud.google.com/architecture/best-practices-for-building-containers">https://cloud.google.com/architecture/best-practices-for-building-containers</a></p>

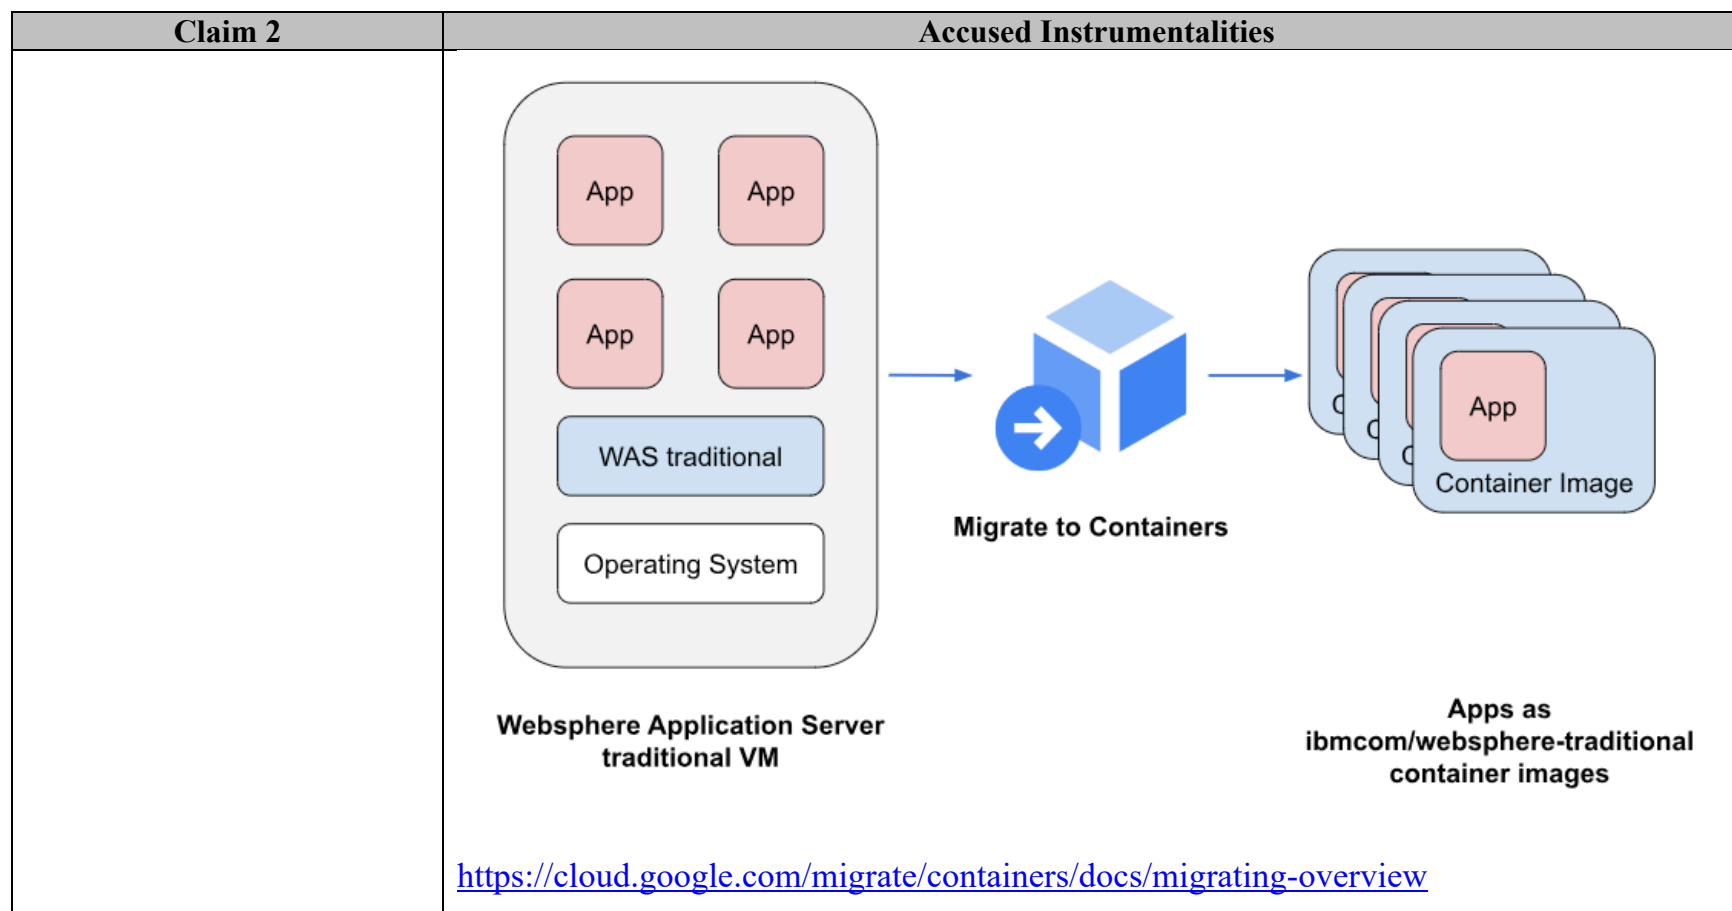
Claim 1	Accused Instrumentalities
	 <p data-bbox="674 885 1060 950"><b>Websphere Application Server traditional VM</b></p> <p data-bbox="1123 690 1396 722"><b>Migrate to Containers</b></p> <p data-bbox="1480 868 1869 966"><b>Apps as ibmcom/websphere-traditional container images</b></p> <p data-bbox="632 1031 1606 1063"><a href="https://cloud.google.com/migrate/containers/docs/migrating-overview">https://cloud.google.com/migrate/containers/docs/migrating-overview</a></p>

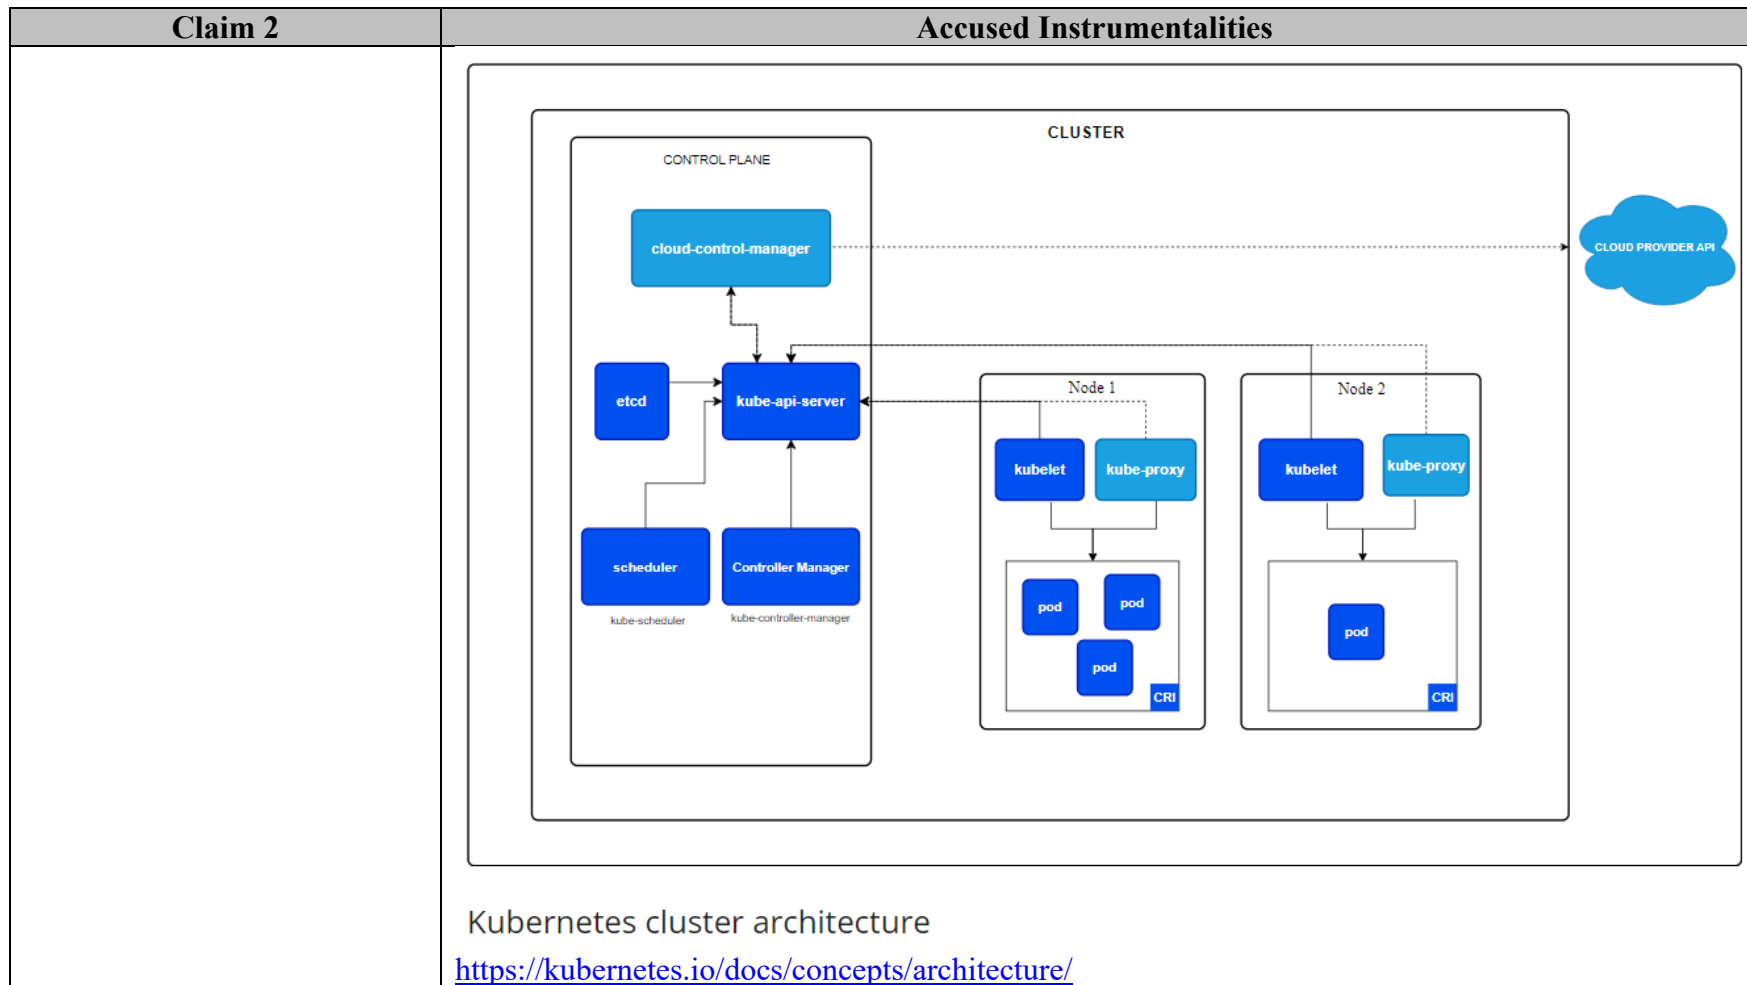
**Claim 2**

Claim 2	Accused Instrumentalities
2. A computing system as defined in claim 1, wherein in operation, multiple instances of an SLCSE stored in the shared library run	Each Accused Instrumentality comprises or constitutes a computing system as defined in claim 1, wherein in operation, multiple instances of an SLCSE stored in the shared library run simultaneously within the operating system.

Claim 2	Accused Instrumentalities
simultaneously within the operating system.	<p>For example, an individual host/node runs multiple containers and/or pods simultaneously, each of which has an instance of an SLCSE.</p> <p><i>See, e.g.:</i></p> <p>Containers solve the portability problem by isolating the application and its dependencies so they can be moved seamlessly between machines. A process running in a container lives isolated from the underlying environment. You control what it can see and what resources it can access. This helps you use resources more efficiently and not worry about the underlying infrastructure.</p> <p>The container image specifies the container's file system. For example, if you're running a Node.js application, the container image would contain your app, Node.js, and other dependencies like Linux system libraries (except the kernel). A container image usually extends a base operating system image, or <b>base image</b>. This base image is the basis of your container, so you'll want to ensure that it's properly patched and free from known vulnerabilities.</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 2	Accused Instrumentalities
	<div data-bbox="699 272 1822 950"> <div> <div>Dockerfile App 1</div> <div>FROM node:19.7.0</div> <div>ADD src_app1 /src/</div> <div>RUN cd /src &amp;&amp; \npm install</div> </div> <div> <div>Dockerfile App 2</div> <div>FROM node:19.7.0</div> <div>ADD src_app2 /src/</div> <div>RUN cd /src &amp;&amp; \npm install</div> </div> <div> <div>Common layers, downloaded only once</div> <div>Layers unique to each image</div> </div> </div> <p data-bbox="632 1015 1677 1052"><a href="https://cloud.google.com/architecture/best-practices-for-building-containers">https://cloud.google.com/architecture/best-practices-for-building-containers</a></p>





Claim 2	Accused Instrumentalities
	<h1 data-bbox="648 261 1125 337">Containers</h1> <p data-bbox="648 399 1906 557">Each container that you run is repeatable; the standardization from having dependencies included means that you get the same behavior wherever you run it.</p> <p data-bbox="648 613 1902 771">Containers decouple applications from the underlying host infrastructure. This makes deployment easier in different cloud or OS environments.</p> <p data-bbox="648 828 1881 985">Each <u>node</u> in a Kubernetes cluster runs the containers that form the <b>Pods</b> assigned to that node. Containers in a Pod are co-located and co-scheduled to run on the same node.</p> <p data-bbox="632 1036 1228 1068"><a href="https://kubernetes.io/docs/concepts/containers/">https://kubernetes.io/docs/concepts/containers/</a></p>

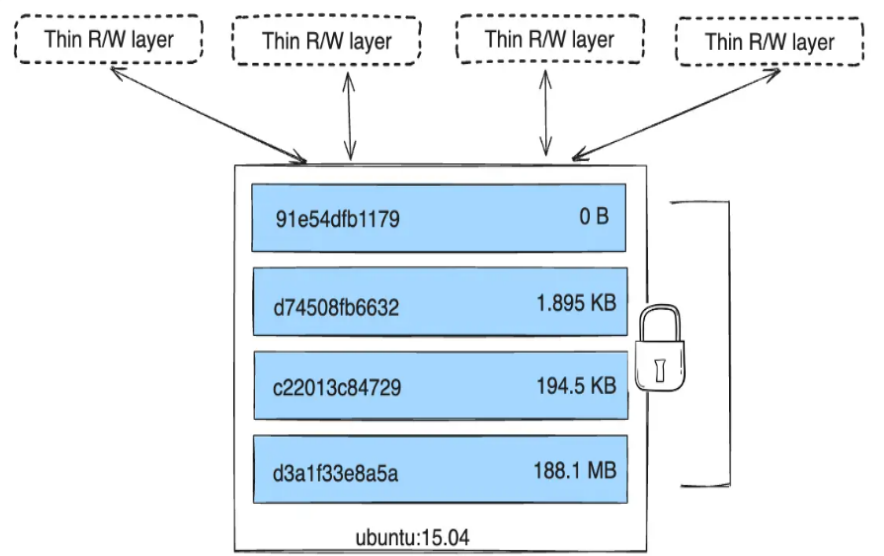


Claim 2	Accused Instrumentalities
	<h1 data-bbox="667 248 1507 321">Kubernetes Scheduler</h1> <p data-bbox="667 370 1640 459">In Kubernetes, <i>scheduling</i> refers to making sure that <u>Pods</u> are matched to <u>Nodes</u> so that <u>Kubelet</u> can run them.</p> <h2 data-bbox="667 581 1314 654">Scheduling overview</h2> <p data-bbox="667 695 1734 938">A scheduler watches for newly created Pods that have no Node assigned. For every Pod that the scheduler discovers, the scheduler becomes responsible for finding the best Node for that Pod to run on. The scheduler reaches this placement decision taking into account the scheduling principles described below.</p> <p data-bbox="667 987 1686 1125">If you want to understand why Pods are placed onto a particular Node, or if you're planning to implement a custom scheduler yourself, this page will help you learn about scheduling.</p> <p data-bbox="634 1174 1554 1206"><a href="https://kubernetes.io/docs/concepts/scheduling-eviction/kube-scheduler/">https://kubernetes.io/docs/concepts/scheduling-eviction/kube-scheduler/</a></p>

Claim 2	Accused Instrumentalities
	<h2 data-bbox="646 250 1209 315">Running containers</h2> <p data-bbox="646 362 1902 483">Docker runs processes in isolated containers. A container is a process which runs on a host. The host may be local or remote. When you execute <code>docker run</code>, the container process that runs is isolated in that it has its own file system, its own networking, and its own isolated process tree separate from the host.</p> <p data-bbox="632 524 1220 557"><a href="https://docs.docker.com/engine/reference/run/">https://docs.docker.com/engine/reference/run/</a></p>

**Claim 3**

Claim 3	Accused Instrumentalities
<p data-bbox="205 732 600 979">3. A computing system according to claim 1 wherein OSCSEs corresponding to and capable of performing the same function as SLCSEs remain in the operating system kernel.</p>	<p data-bbox="632 732 1902 834">Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein OSCSEs corresponding to and capable of performing the same function as SLCSEs remain in the operating system kernel.</p> <p data-bbox="632 862 1856 964">For example, both Docker and Kubernetes systems preserve the host kernel substantially unchanged; therefore the OSCSEs corresponding to the SLCSEs remain in the operating system kernel.</p> <p data-bbox="632 992 747 1024"><i>See, e.g.:</i></p> <p data-bbox="638 1060 1430 1187">Most base images are basic or minimal Linux distributions: Debian, Ubuntu, Redhat, Centos, or Alpine. Developers usually consume these images directly from Docker Hub, or other sources. There are official providers along with a wide variety of other downstream repackagers that layer software to meet customer needs.</p> <p data-bbox="632 1192 1587 1224"><a href="https://cloud.google.com/software-supply-chain-security/docs/base-images">https://cloud.google.com/software-supply-chain-security/docs/base-images</a></p>

Claim 3	Accused Instrumentalities
	<p><a href="#">Container image files</a> are complete, static and executable versions of an application or service and differ from one technology to another. <a href="#">Docker images</a> are made up of multiple layers, which start with a base image that includes all of the dependencies needed to execute code in a container. Each image has a readable/writable layer on top of static unchanging layers. Because each container has its own specific container layer that customizes that specific container, underlying image layers can be saved and reused in multiple containers. An Open Container Initiative (OCI)</p> <p><a href="https://www.techtarget.com/searchitoperations/definition/container-containerization-or-container-based-virtualization">https://www.techtarget.com/searchitoperations/definition/container-containerization-or-container-based-virtualization</a></p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

**Claim 4**

Claim 4	Accused Instrumentalities
<p>4. A computing system according to claim 1 wherein the one or more SLCSEs provided to one of the plurality of software applications having exclusive use thereof, use system calls to access services in the operating system kernel.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein the one or more SLCSEs provided to one of the plurality of software applications having exclusive use thereof, use system calls to access services in the operating system kernel.</p> <p>For example, the SLCSEs in a container use system calls to access services in the operating system kernel. For example, the glibc library (or other similar library) in the container uses system calls to interface with the host Linux kernel. In general, system calls can be observed using a tool such as strace.</p> <p><i>See, e.g.:</i></p> <p>The <b>GNU C Library</b>, commonly known as <b>glibc</b>, is the <a href="#">GNU Project</a> implementation of the <a href="#">C standard library</a>. It is a wrapper around the system calls of the <a href="#">Linux kernel</a> for application use. Despite its name, it now also directly supports <a href="#">C++</a> (and, indirectly, other <a href="#">programming languages</a>). It was started in the 1980s by the <a href="#">Free Software Foundation</a> (FSF) for the <a href="#">GNU</a> operating system.</p> <p><a href="https://en.wikipedia.org/wiki/Glibc">https://en.wikipedia.org/wiki/Glibc</a></p>

We can now get the process id directly from the cgroup. It will be located in the cgroup.procs file.

### Terminal 2 - Worker Node ###

# Get the process id

```
$ cat cri-containerd-ceeeef06afe89c8223d33b11e8d9e0b207118ac4dac3af826687668ee1ee
16254
```

# Validate what is running under the process

```
$ ps aux | grep 16254
```

```
azureus+ 16254 0.0 0.1 713972 10476 ? Ssl 15:04 0:00 ./faultyapp
azureus+ 94806 0.0 0.0 7004 2168 pts/0 S+ 16:22 0:00 grep --color=a
```

Got it! With that, we can try to find out what is going out inside the app. Lets try to run strace to get some more insight.

### Terminal 2 - Worker Node ###

```
$ sudo strace -p 16254 -f
```

```
...
```

# The app is trying to read a file port.txt

```
[pid 16269] openat(AT_FDCWD, "port.txt", O_RDONLY|O_CLOEXEC <unfinished ...>
```

```
[pid 16254] epoll_pwait(5, <unfinished ...>
```

# The file does not exist

```
[pid 16269] <... openat resumed> = -1 ENOENT (No such file or directory)
```

```
[pid 16254] <... epoll_pwait resumed>[], 128, 0, NULL, 0) = 0
```

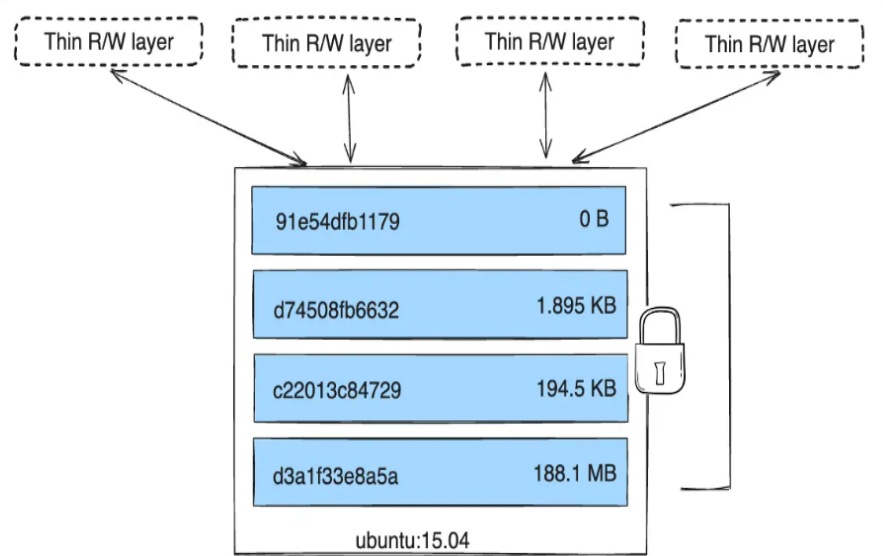
```
[pid 16269] write(1, "Something went wrong...\n", 24 <unfinished ...>
```

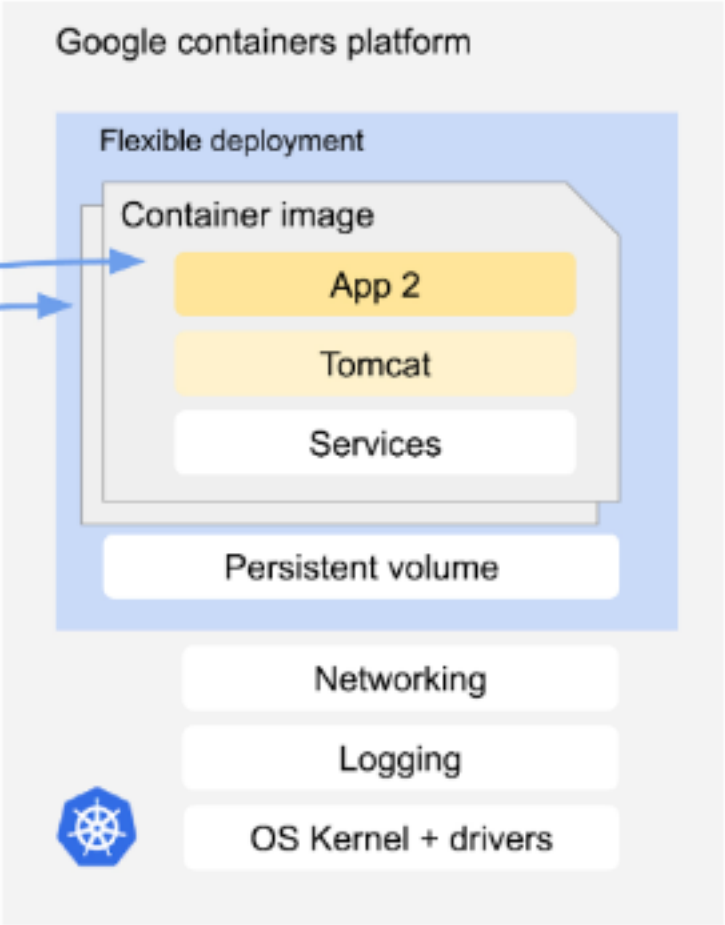
After filtering the output, we can see the application is trying to read a text file called port.txt, and a few lines later, there is a message stating ENOENT (No such file or directory). Let's create that file.

<https://www.berops.com/blog/a-different-method-to-debug-kubernetes-pods>

**Claim 5**

Claim 5	Accused Instrumentalities
<p>5. A computing system according to claim 1 wherein the operating system kernel comprises a kernel module adapted to serve as an interface between an SLCSE in the context of an application program and a device driver.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein the operating system kernel comprises a kernel module adapted to serve as an interface between an SLCSE in the context of an application program and a device driver.</p> <p>For example, the server (node) includes an operating system having a kernel. The kernel comprises a kernel module which enables applications (including their libraries) to have access to system resources such as storage, <i>i.e.</i>, acts as an interface between applications/libraries and OS libraries or drivers</p> <p><i>See, e.g.:</i></p> <p><b>Container images</b></p> <p>A <a href="#">container image</a> is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p><a href="https://kubernetes.io/docs/concepts/containers/">https://kubernetes.io/docs/concepts/containers/</a></p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p><a href="https://www.techtarget.com/searchitoperations/definition/Docker-image">https://www.techtarget.com/searchitoperations/definition/Docker-image</a></p>

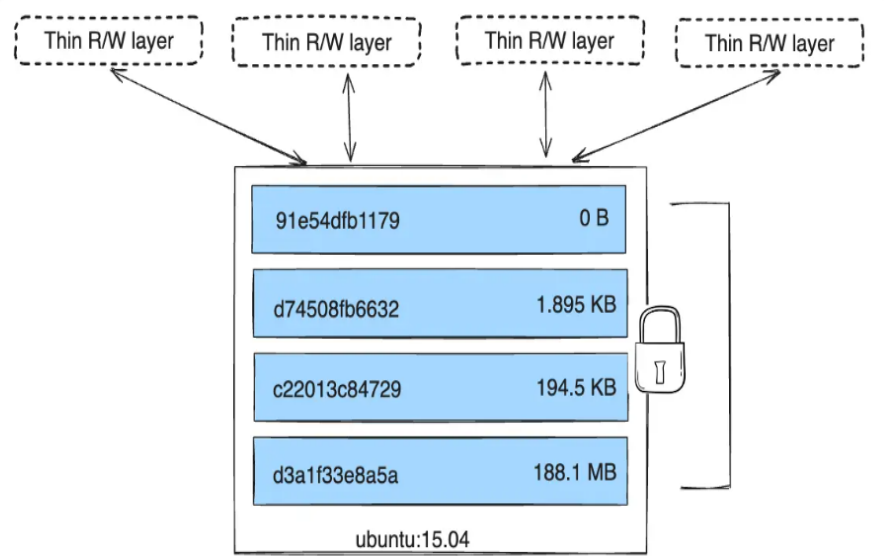
Claim 5	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p> <p>Containers use specific features of the Linux kernel that “trick” individual applications into thinking they’re in their own unique environment, even though multiple applications share the same host kernel. (If you’re not familiar with the Linux kernel, it’s a part of the operating system that communicates between processes--requests that do user tasks like opening a file, running a program-- and the hardware. It manages resources like memory and CPU to meet these requests).</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

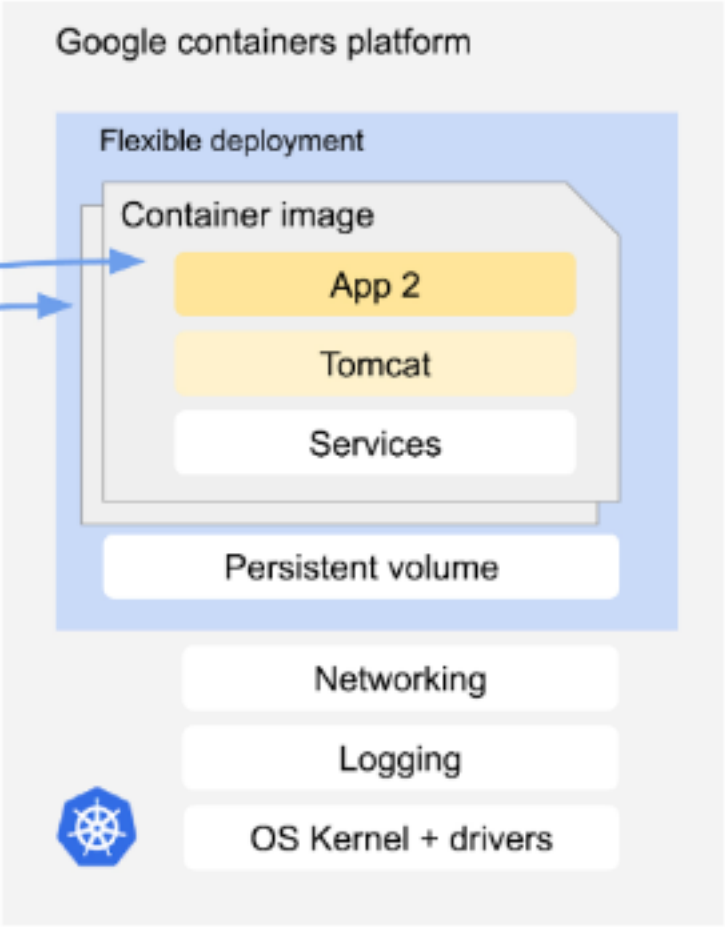
Claim 5	Accused Instrumentalities
	 <p>The diagram illustrates the Google containers platform architecture. At the top is the 'Google containers platform' layer. Below it is the 'Flexible deployment' layer, which contains a 'Container image' box. This box is divided into three sections: 'App 2' (yellow), 'Tomcat' (yellow), and 'Services' (white). Two blue arrows point from the left towards the 'Container image' box. Below the 'Container image' box is a 'Persistent volume' box. Below the 'Flexible deployment' layer are three stacked boxes: 'Networking', 'Logging', and 'OS Kernel + drivers'. A Docker logo is positioned to the left of the 'OS Kernel + drivers' box.</p> <p><a href="https://cloud.google.com/blog/products/application-modernization/shift-your-apps-to-container-based-workloads-on-the-command-line">https://cloud.google.com/blog/products/application-modernization/shift-your-apps-to-container-based-workloads-on-the-command-line</a></p>



**Claim 10**

Claim 10	Accused Instrumentalities
<p>10. A computing system according to claim 2 wherein SLCSEs stored in the shared library are linked to particular software applications of the plurality of software applications as the particular software applications are loaded such that the particular software applications have a link that provides unique access to a unique instance of a CSE.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system according to claim 2 wherein SLCSEs stored in the shared library are linked to particular software applications of the plurality of software applications as the particular software applications are loaded such that the particular software applications have a link that provides unique access to a unique instance of a CSE.</p> <p>For example, the containers can share common dependencies and components using layered images, and multiple containers can use the same base image. Therefore, each container, containing the application software running under the operating system of the server hosting GKE, uses a unique instance of the corresponding critical system element to execute the respective application software and has a link to that unique instance.</p> <p><i>See, e.g.:</i></p> <h2 style="text-align: center;">Container images</h2> <p>A <a href="#">container image</a> is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p><a href="https://kubernetes.io/docs/concepts/containers/">https://kubernetes.io/docs/concepts/containers/</a></p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p><a href="https://www.techtarget.com/searchitoperations/definition/Docker-image">https://www.techtarget.com/searchitoperations/definition/Docker-image</a></p>

Claim 10	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p> <p>Containers use specific features of the Linux kernel that “trick” individual applications into thinking they’re in their own unique environment, even though multiple applications share the same host kernel. (If you’re not familiar with the Linux kernel, it’s a part of the operating system that communicates between processes--requests that do user tasks like opening a file, running a program-- and the hardware. It manages resources like memory and CPU to meet these requests).</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 10	Accused Instrumentalities
	 <p>The diagram illustrates the Google containers platform architecture. At the top is the 'Google containers platform' layer. Below it is the 'Flexible deployment' layer, which contains a 'Container image' box. This box is divided into three sections: 'App 2' (yellow), 'Tomcat' (yellow), and 'Services' (white). Two blue arrows point from the left towards the 'Container image' box. Below the 'Container image' box is a 'Persistent volume' box. Below the 'Persistent volume' box are three stacked boxes: 'Networking', 'Logging', and 'OS Kernel + drivers'. A Docker logo is positioned to the left of the 'OS Kernel + drivers' box.</p> <p><a href="https://cloud.google.com/blog/products/application-modernization/shift-your-apps-to-container-based-workloads-on-the-command-line">https://cloud.google.com/blog/products/application-modernization/shift-your-apps-to-container-based-workloads-on-the-command-line</a></p>

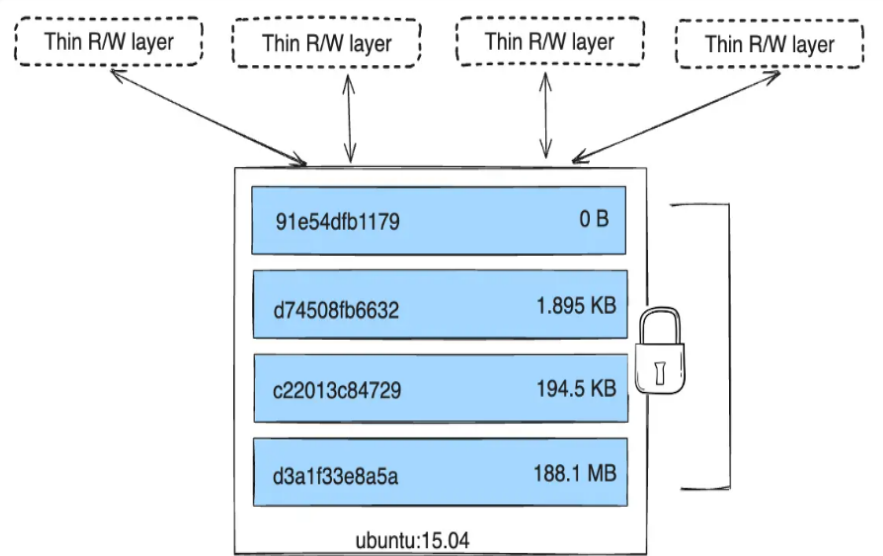
Claim 10	Accused Instrumentalities
	<p>Containers solve the portability problem by isolating the application and its dependencies so they can be moved seamlessly between machines. A process running in a container lives isolated from the underlying environment. You control what it can see and what resources it can access. This helps you use resources more efficiently and not worry about the underlying infrastructure.</p> <p>The container image specifies the container's file system. For example, if you're running a Node.js application, the container image would contain your app, Node.js, and other dependencies like Linux system libraries (except the kernel). A container image usually extends a base operating system image, or <b>base image</b>. This base image is the basis of your container, so you'll want to ensure that it's properly patched and free from known vulnerabilities.</p> <p><a href="https://services.google.com/fh/files/misc/why_container_security_matters.pdf">https://services.google.com/fh/files/misc/why_container_security_matters.pdf</a></p>

Claim 10	Accused Instrumentalities
	<div style="display: flex; justify-content: space-around; align-items: flex-start;"> <div style="text-align: center;"> <p>Dockerfile App 1</p> <pre style="background-color: #f8d7da; padding: 10px; margin-bottom: 5px;">FROM node:19.7.0</pre> <pre style="background-color: #d4edda; padding: 10px; margin-bottom: 5px;">ADD src_app1 /src/</pre> <pre style="background-color: #d4edda; padding: 10px;">RUN cd /src &amp;&amp; \   npm install</pre> </div> <div style="text-align: center;"> <p>Dockerfile App 2</p> <pre style="background-color: #f8d7da; padding: 10px; margin-bottom: 5px;">FROM node:19.7.0</pre> <pre style="background-color: #d4edda; padding: 10px; margin-bottom: 5px;">ADD src_app2 /src/</pre> <pre style="background-color: #d4edda; padding: 10px;">RUN cd /src &amp;&amp; \   npm install</pre> </div> </div> <div style="margin-top: 20px;"> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="width: 20px; height: 15px; background-color: #f8d7da; margin-right: 5px;"></div> <span>Common layers, downloaded only once</span> </div> <div style="display: flex; align-items: center;"> <div style="width: 20px; height: 15px; background-color: #d4edda; margin-right: 5px;"></div> <span>Layers unique to each image</span> </div> </div> <p style="color: blue; text-decoration: underline;"> <a href="https://cloud.google.com/architecture/best-practices-for-building-containers">https://cloud.google.com/architecture/best-practices-for-building-containers</a> </p>

**Claim 18**

Claim 18	Accused Instrumentalities
18. A computer system as defined in claim 2 wherein SLCSEs are not copies of OSCSEs.	<p>Each Accused Instrumentality comprises or constitutes a computer system as defined in claim 2 wherein SLCSEs are not copies of OSCSEs.</p> <p>For example, in a typical case the SLCSEs come from a Linux distribution independent of the host operating system, and thus are not identical to the OSCSEs. For another example, the SLCSEs are</p>

Claim 18	Accused Instrumentalities
	<p>provided to the computer system through a separate process than the process by which the OSCSEs are provided to the computer system, and thus are not copied from the OSCSEs.</p> <p><i>See, e.g.:</i></p> <h2 data-bbox="674 402 1062 456">Container images</h2> <p data-bbox="674 492 1417 641">A <a href="#">container image</a> is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p data-bbox="632 667 1228 699"><a href="https://kubernetes.io/docs/concepts/containers/">https://kubernetes.io/docs/concepts/containers/</a></p> <p data-bbox="646 755 1696 917">Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p data-bbox="632 943 1535 976"><a href="https://www.techtarget.com/searchitoperations/definition/Docker-image">https://www.techtarget.com/searchitoperations/definition/Docker-image</a></p>

Claim 18	Accused Instrumentalities
	<p data-bbox="646 245 1898 375">Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p data-bbox="632 1019 1224 1052"><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>